

Culpable ignorance in a collective setting

Forthcoming in *Acta Philosophica Fennica* 94 (2018), pp. 7-34.

Säde Hormio, Practical Philosophy, University of Helsinki

Abstract

This paper explores types of organisational ignorance and ways in which organisational practices can affect the knowledge we have about the causes and effects of our actions. I will argue that because knowledge and information are not evenly distributed within an organisation, sometimes organisational design alone can create individual ignorance. I will also show that sometimes the act that creates conditions for culpable ignorance takes place at the collective level. This suggests that quality of will of an agent is not necessary to explain culpable ignorance in an organisational setting.

1. Introduction

Ignorance is traditionally seen as an excuse for blame. While there has been a growing interest in ignorance and individual responsibility,¹ literature on collective ignorance and what it means for responsibility has so far been quite thin on the ground. This paper hopes to bridge some of the gaps in the literature by exploring ways in which organisational practices can affect the knowledge we have about the causes and effects of our actions.

My concern is with the *epistemic condition of responsibility*, of acting under ignorance, not with *epistemic responsibility* (i.e. questions regarding what you ought or ought not to believe). If one does not know what one is involved in, and cannot be reasonably expected to know either, then one cannot be praiseworthy or blameworthy. But what about the collective? Organisational practices can affect the knowledge we have about the causes and effects of our actions. Can an organization be blameworthy when an individual acts under ignorance?

There are three interlinked issues that I will explore in this paper. The first is about discovering what different types of ignorance can be found in organisations. The second describes how sometimes organisational design creates ignorance even without anyone trying deliberately to mislead anyone. The third concerns how organisations can be responsible for an individual's culpable ignorance. I explore these questions mainly through fragmentation of information and suppression of knowledge.

¹ See for example the collection of papers in Peels 2016 and Robichaud and Wieland 2017.

I will discuss responsibility at two levels: at the level of the organisation and at the level of individual members of the organisation. Organisational ignorance can be ignorance about facts or it can arise from the suppression of knowledge. My discussion of members in an organisation is focused on the “regular employees” in a large organization who are subjected to the standard effects of compartmentalization of information that is necessary in large organised collectives. The assessment of excusing or culpability of individual ignorance is therefore in most cases aimed at this “general” level, not at individuals that occupy special positions in the organisation with privileged access to information, such as directors. The assessment of their individual responsibility might vary a lot from the responsibility of the average employee, depending on circumstances.

This paper proceeds as follows. I begin by discussing ignorance from the point of view of the collective. Section 2 introduces types of ignorance about knowledge and facts, while section 3 discusses ignorance arising from suppression of knowledge. Section 4 looks at organisational design and how fragmentation of information can lead to ignorance. I then turn my attention to individual members of collectives in section 5, where I discuss culpable ignorance in collective setting and suggest that in some cases, the blame for culpable ignorance should be directed at the level of the organisation.

2. Ignorance about knowledge and facts

The role of ignorance in society is not just passive and negative. Often ignorance is simply unavoidable or neutral. No-one can know everything, and not all information is relevant for all. Ignorance is also an indispensable element in many social relations and structures (Moore and Tumin 1949). Ignorance is not a static state of affairs. In the course of time, unknowns are transformed into new knowledge, while at the organisational level some old knowledge is forgotten and replaced with ignorance (Roberts 2013, p. 218). Science challenges the existing body of knowledge and is at the centre of most of our efforts of turning unknowns and ignorance into knowledge.

This section introduces issues related to ignorance about knowledge and facts, while in the next section ignorance arising from the suppression of knowledge is discussed. I will use Joanne Roberts’s (2013) work on organisational ignorance as my starting point.² While the concern in Roberts’s original work is

² Roberts (2013) divides ignorance into three key sources: ignorance arising from the absence of knowledge, ignorance about knowledge, and ignorance arising from the suppression of knowledge. This paper concentrates on the latter two, but I will briefly describe what is involved in the first category before setting it aside. *Known unknowns* are things outside the limits of

not with normative questions or moral responsibility, I will apply her categories in this way.³

Ignorance about knowledge and facts can be divided into three subcategories. *Knowable recognised unknowns* are something that could be found out given the right resources and motivation. In Roberts's paper these recognised unknowns that are knowable (i.e. the ignorance about the facts could be overcome) are called "knowable known unknowns", but this term can be philosophically confusing due to the tautology involved. Therefore I am referring to this category as knowable recognised unknowns throughout. This ignorance could be overcome if the organisation wanted to, i.e. there is no missing science or technology standing in the way of obtaining knowledge about these facts.⁴ Knowable recognised unknowns are either outside the focus of the organisation, resulting in lack of motivation to overcome the ignorance, or the costs and benefits do not add up to put the expenditure necessary into resources to overcome it. To give an example, an educational organisation working in India should be sensitive to issues to do with discrimination against Dalit students, while a similar organisation working in Argentina does not necessarily need to know anything about the issue.

Unknown knowns are things we do not know that we know, including tacit knowledge, meaning that this type of ignorance does not necessarily prevent the use of the knowledge. In fact it may underpin creativity in the form of intuition. In organizations unrecognized knowledge is often embedded in

our knowledge, things we know that we do not know. *Unknown unknowns* are beyond anticipation, a total lack of knowledge: something we are not even aware of being ignorant about at a specific point in time, so they cannot be directly investigated. Examples of both can be found in astrophysics and the current knowledge of our solar system. Known unknowns refers to a state of ignorance at a specific point in time in an organizational context, it is an awareness that certain knowledge is not in fact known by the organisation or its members. Organizational known unknowns can drive research and development, leading to innovations. They are also something that needs to be taken into account when the organisation is involved with research and development that carries high risks. (Roberts 2013, p. 217-221). Due to this, there could be some cases of collective responsibility linked to known unknowns, but even in them the potential blame would not be likely linked to culpable ignorance. When it comes to unknown unknowns, undoubtedly advances in science will bring new knowledge that is currently outside the scope of our knowledge (and create new unknowns in the process), or even beyond our anticipation. However, as unknown unknowns are so completely outside our control, I do not find the type to be relevant for debates on blameworthy ignorance.

³ I make no claim to offer an exhaustive list of all possible types of ignorance in an organisational setting, but I believe that Roberts's work offers a comprehensive typology. Management science is an interdisciplinary field where the goal is to find solutions to organizational problems and challenges, and her work draws from research into ignorance from a variety of fields, including sociology, politics, economics and philosophy. I therefore believe it to be a fertile ground to start mapping out normative framework on ignorance.

⁴ In contrast to known unknowns mentioned in footnote 2.

routines and collective practices, existence of which is usually exposed only in retrospect once it is lost, like when people retire. Finally *errors* arise from inaccuracy, confusion, uncertainty or incompleteness. They are the things we think we know, but don't. The more complex an organization is, the more it is prone to errors. Individuals might either wrongly assess their level of competence, resulting in an error that has implications for the wider organisation, or there is a system failure due to a design fault in the organization. Changing environments increase the risk of organizational errors occurring. (Roberts 2013, p. 218-223).

I will soon discuss how knowable recognised unknowns and errors can raise issues of responsibility and blameworthiness in a collective setting. However, I will first argue that it is harder to see how unknown knowns could raise these issues. Tacit knowledge is difficult to articulate and sharing it is not a simple matter, especially in an organizational setting. While one could argue that it should be part of good managerial practise to try to ensure that as much as possible of the silent knowledge is passed on to the new employees as old ones retire, for example, the fact that it is hard to measure when this passing on has been successful (or not) belies that tacit knowledge is not a very fruitful basis for blame arising from ignorance. Of course, if an organisation makes these kind of practices difficult or impossible (for example, letting go the old staff before new staff starts, not asking for any handover notes etc.), it could be argued to be not just engaged in deficient investigation, but also involved in preventing subsequent discovery (categories that I will come to in section 5). Still, I find that genuine cases of culpable ignorance remain limited for this type of ignorance.

To put knowable recognised unknowns and errors in more concrete terms, I will introduce the tale of the exploding toasters. The example will help to illustrate responsibility issues related to ignorance in collective contexts, in this narrative in supply chain management. Supply chain management is a term with many usages, as it covers the product cycle from design of new products and services to the delivery of the finished product to the end customers (Lu and Swaminathan 2015). I will use the term here to refer to procurement of goods and to how companies manage outsourcing of the manufacturing of their products.

Violet works as an in-byer for a large retail chain Sell-A-Lot with shops all over the world. She finalises a large toaster purchase order. Unbeknownst to Violet, the toaster company, Exciting Electronics, has very recently changed their manufacturer to cut costs. The decision to make changes to the supply chain was made at the last minute with regards to the upcoming peak sales period. As a consequence, the new toasters are manufactured in a rush by the new supplier and some have loose wires. Exciting

Electronics sales representative Sharon is not aware of the problem and sells faulty toasters to Sell-A-Lot in good faith. Before long shop managers at Sell-A-Lot are flooded with angry phone calls from customers about their brand-new toasters sending off sparks. One unlucky soul has her house set on fire but lives to see another day. There need not exist any malicious intent; rather, we are looking at negligence. Nobody in the new factory supplying Exciting Electronics makes the wires loose on purpose and with evil intent, all faults are due to the too hectic manufacturing process.

Errors can raise questions of blameworthiness in many ways, the most serious ones being systematic errors. A systematic error of some sort must have taken place at Exciting Electronics for Sharon to be able to sell faulty toasters to a customer in the first place. It looks like a clear case of design fault in the organization if sales are allowed to go through on faulty products if the company knows about the fault. Systemic errors are blameworthy, as I will argue later on. In case of genuine errors, ones that are not due to negligent practises, there will be no blame (although there might be certain sense of agent-regret as I argue in section 5).

Knowable recognised unknowns were something that could be found out given the right resources and motivation. With this category, I think that we should keep in mind that organisations have less excuses of being ignorant of facts than individuals do. They can set aside money to fund research into a suitable course of action for them and have a group of experts dedicate their working hours to thinking through an issue from the organisation's point of view. If the required expertise cannot be found among their existing members, they can hire new staff or employ consultants. They can and should do this when new issues arise that affect their operating environment and future operations. Saying that, if something is defensibly outside the focus of the organisation, leading to lack of motivation to overcome the ignorance, I find that the resulting ignorance is not susceptible to blame. In the case of Exciting Electronics, this could be the environmental impact of hairdryers, for example, if they did not manufacture any such products.

Trickier variety of knowable recognised unknowns are cases where the costs and benefits do not justify the expenditure to attempt to overcome ignorance, so the decision to remain ignorant is made on a purely financial basis. I say trickier because cost and benefit analyses of any sort are riddled with normative assumptions. One board of directors in one company might deem something to be too costly, while another one would rule it to be a justifiable expenditure. If Sell-A-Lot or Exciting Electronics has decided that it is too costly to put in the resources necessary to adequately monitor their supply chain in relation to, say, the working conditions in the factories, then these working

conditions remain knowable recognised unknowns to the organisations. Arguably this situation could (and should) be different, but such arguments need to be rooted in normative considerations outside simple cost and benefit analyses.

More broadly put, when it comes to knowable recognised unknowns, the issue of whether ignorance is blameworthy or excusing can only be settled in the context of deciding the adequate focus for the organisation. Following Tuomela (2007, p. 15) and Laitinen (2014, p.218), this focus could be labelled as the *ethos* of the organisation, consisting of the central questions and practical matters that are vital to the purpose of the group (the group's *realm of concern*) and the answers it has collectively accepted to be its view (*intentional horizon*). Ethos thus covers the central goals and commitments of the organisation. The exploding toasters would be a case of knowable recognised unknowns if Sell-A-Lot has poor supply chain management and little motivation to invest properly in even basic product safety testing, let alone other corporate social responsibility measures. The harm caused by the faulty toasters could then be traced back to corporate policies and priorities in addition to the obvious fault on the manufacturing side. In other words, the ethos of the organisation could be argued to include negligence towards safety.

What is important for the topic at hand is that while the ethos determines a group's identity (and is part of what marks its continuation together with its historical and modal properties), it is in a state of flux (to what degree varies naturally a lot between each case). The ethos of a group is therefore not set in stone, as elements of it may change (and almost always do to some degree at least, especially when it comes to the large organisational-size groups). Examples are everywhere: corporations venture into new areas of production, political parties amalgamate new goals, the jurisdiction of local authorities change, university begins to offer courses in a new subject matter; it is easy to keep coming up with examples. Therefore, to simply state that some knowledge is currently outside the realm of concern of the ethos of an organisation does not, by itself, settle much in terms of responsibility and possible blameworthiness.

Because organisations are responsive to their environments and must regularly review their realm of concerns as well their intentional horizons, I argue that knowable recognised unknowns are always a normative matter. They could provide an interesting angle for political philosophy to look at questions related to what should fall within the ethos of the organisation. For example, how much and to what degree should a government be aware of the impacts of the actions of the banks it has chartered? Or looking at corporate responsibility, how far into their supply chain should a corporation look?

However, ignorance about knowledge and facts is not the only category of organisational ignorance that is ripe for philosophical analysis on responsibility. I turn to ignorance that results from suppression of knowledge in the next section.

Before moving on, I will briefly address the question of what organisations can be said to know. I have so far discussed knowledge that the organisation does not have, but what about the knowledge an organisation can be argued to have? Imagine a spy ring of some kind, where the spies do not know the identity of the other spies or have access to the information the others have.⁵

In my example, each spy has been assigned a code name and a secret phone with which to get in touch with the others. They have an assignment to complete where Aja knows the target, Katya the method, and Shea the time and the place. The person who set them up on the mission was involved in a car crash and is lies in a coma in hospital somewhere. Shea has been instructed to invite the other two to come to the designated place at the designated time, Katya to bring the means, and Aja to put it in use towards the target. Between them, they have all the information necessary to successfully complete the assignment of the spy ring. However, can the spy ring as a collective be said to have knowledge about what the assignment is (before it is carried out)?

I suggest that the spy ring does have working knowledge about the assignment, as it is able to carry it out. If the spy ring operates as part of some organised espionage group, it is the institution rather than the spy ring per se that knows what the assignment is.⁶ However, if the person who is now in coma is some rogue agent (or some eccentric millionaire who engages in espionage as a hobby etc.), then the spy ring per se has the knowledge. Saying that, the knowledge of the spy ring is highly fragmented. Therefore the group knowledge is not robust at all: if one link was missing, they would not be able to achieve their goal. In this case, the spy ring knows *how* to perform the assignment, but not *that* the assignment is to X.⁷ After they have carried it out, the spy ring understandably also knows more about the assignment - who was the target, where, and the means - as the knowledge of individual group members has come together in action. Knowing both how to do something and the details of what you are doing is clearly more robust group knowledge than knowing just the former.

⁵ I thank Jaakko Kuorikoski for suggesting using a spy ring as an example.

⁶ I thank Deborah Tollefsen for pushing me on this point.

⁷ I am indebted to an anonymous reviewer for making me draw up further distinctions on this point.

By bringing in robustness, I am suggesting that group knowledge comes in degrees. Furthermore, I suggest that making the group knowledge more robust could in some cases be thought of as something that the collective should do. With highly fragmented information, the collective is taking a risk that the goal is not reached. In the case of spy rings, this risk seems acceptable, as there are benefits to the arrangement: if Katya is captured, the secret mission is not revealed. However, in many everyday cases the information a collective has could be too fragmented and the associated risks not acceptable, like with the loose wires and internal practices of Exciting Electronics. In these cases, the ethos of the collective could be argued to be too negligent towards supply chain management. I will return to these issues in section 4, where I discuss organisational fragmentation of information.

3. Suppression of knowledge

This section introduces types of organisational ignorance that can arise from the suppression of knowledge. Withholding some important information, or the tendency to only communicate the positive news, is common among corporations and other large modern organisations. This section looks at different reasons for suppressing knowledge.

Taboos are socially constructed bans on certain types of knowledge deemed to be polluting. They can also be actively cultivated within organisations to influence the way its members behave, like a taboo about discussing bullying in the workplace. When knowledge is too painful to acknowledge, or it does not fit with one's worldview, it can be repressed or ignored, resulting in *denials*. Organisational denials can lead to ignoring evidence that contradicts the group decision for the sake of unanimity. This can be especially dangerous when encouraged by those in charge, as toleration for recklessness and dishonesty in practices has a tendency to spread. Denials can also be used strategically, like when a company encourages ignorance in their customers through misinformation campaigns. We talk of *secrecy* when knowledge is consciously suppressed by individuals or collectives. Pockets of ignorance can be deliberately created for power purposes. Some secrecy is essential (keeping trade secrets, for example) but there has to be a balance and an understanding of how much secrecy the stakeholders are willing to tolerate. *Privacy* is socially sanctioned secrecy and the right to privacy is enshrined in many laws and declarations. To build trust between an organisation and its members and stakeholders, it is important to recognise and protect privacy, for example, the customer data registry of a company. (Roberts 2013, p. 218-226).

I find that privacy could also be thought as a sub-category of secrecy. Privacy here refers to the privacy of the organisations' employees or customers. It is often about trust, the disability of an employee need

not become common knowledge within the workplace, and not keeping customer data safe can by itself be a morally blameworthy act that can lead to the loss of those customers, or even to harm for those customers, depending on your line of business. Organisational ignorance always has power dimensions and the potential to be political in nature. To give an example, privacy has been in the news a lot lately with investigations into how Facebook has handled its users' data. Still, I will set this category aside for the rest of the paper, as my concern is with cases where organisations can make or keep individuals ignorant about issues they arguably should not be, not with issues to do with mismanaging data.

It seems clear that many instances of ignorance arising from the other forms of suppression of knowledge can be amenable to blameworthiness. Think of some corporation that is involved in practices that are questionable, but is not communicating this to its employees or other stakeholders. A recent real-life example of this is how many employees of Google were taken aback when they found out that the company is involved developing algorithms for military use. Google's participation in Project Maven with the Pentagon prompted thousands of its employees to sign an open letter urging Google to not be involved in developing military Artificial Intelligence. When it comes to suppression of knowledge, an organisation has introduced a condition (through denial, secrecy, or taboo) – or failed to remove it – which made it difficult for employees to acquire true belief about the wrongness of being involved in some particular collective action. Depending on the actual circumstances, the employees' ignorance could be culpable if it is due to deficient inference, or excusable if the organisational barrier for acquiring the knowledge is too high, but I will return to this section 5.

Secrecy is very common in organisations. Although we live in the information age, secrecy is still a very prevalent component of our societies. For example, Galison (2008, p. 38) describes how “we are living in a climate of augmented secrecy” today, with the number of classified document pages outnumbering the amount of open literature entering the public libraries and archives each year in the U.S.⁸ Secrecy naturally has a large role to play in organisations such as intelligence services, but it is an ubiquitous part of the corporate world also. Secrecy has positive features for an organisation, like affording more freedom in negotiating difficult situations in politics, or giving a competitive advantage to a corporation by ensuring first-mover advantage in a new product area (Dufresne and Offstein 2008). Still, secrecy can also be used as an excuse. Coming back to the example I have been using, if Exciting Electronics had a policy in place where all information about suppliers were classified as trade secrets, it would be impossible for Sell-A-Lot to know where the products it sells originate from. If the ethos of Sell-A-Lot

⁸ Galison (2008, pp. 37-39) attributes this rise of modern censorship mainly to the infrastructure created after the Second World War around nuclear science and intelligence services.

includes ensuring safe working conditions in its supply chain, this kind of non-transparency from the part of Exciting Electronics would be unacceptable for them, resulting in them procuring their toasters elsewhere.

While genuine trade secrets are one thing, like Google's algorithm or the closely guarded recipe of Coca-Cola, corporations can use secrecy simply to avoid awkward questions about their products or their supply chain without acceptable reasons for their secrecy. The acceptability has to be linked to the kind of reasons given for the secrecy: would they stand the scrutiny of objective outsiders? If there are no acceptable reasons for secrecy, I find that non-transparent business practices are not justifiable, at least in the modern world, where concerns about the treatment of workers in global supply chains or the environmental impacts of rampant consumerism have become widely known.

The problem with organisational secrecy is of course that often we do not know what we should be concerned about, at least until a whistle-blower alerts us to the facts. Still, there are areas that are widely known to be riddled with problems, like the conditions under which many of our clothes or consumer electronics are produced under. Therefore, unless the corporation is forthcoming about the way it handles the problematic areas in its supply chain, there is a high likelihood that they have a thing or two to hide. To use Exciting Electronics again, it could be the case that the top management were aware of the faulty toasters, but decided to deny this to protect profits in the short term. Maybe the quantities already delivered to foreign retailers were large, and costs of recalling products were deemed too great, because faults were found in only a few toasters so far, so they decided to gamble. Let's say Bianca is a conscientious middle-level manager supervising Sharon's department. She was aware that their toaster manufacturer was changed and wanted to make sure that the new products were tested for safety. Bianca enquired after the results from the top managements, who denied knowledge of any faults. Bianca gave Sharon's team the go ahead to sell the toasters to Sell-A-Lot and other domestic retailers. Any blame should be directed at the top management in this case, or the organisation itself, depending on the details.

Denials, on the other hand, can take form of flat-out lies or more subtle agnotology, where the goal is to create misinformation (Proctor and Schiebinger 2008). I have argued elsewhere that an organisation can be blameworthy for distorting public debate through strategic denials (Hormio 2017), but here I want to address only ignorance within organisations, so I will give an example of internal denials. The loose wires could bring back painful issues for Exciting Electronics, so the information could be repressed or ignored by those who are privy to the manufacturing problems. Say that Exciting

Electronics used to own its own factories, but changed its business models some years back in order to become more competitive with its prices. As part of this process, they closed down their own factories in Europe and outsourced their manufacturing to countries with cheap labour and laxer regulations. This decision was far from unanimous at the board level and raised a lot of debate at other levels of the organisation too, let alone among those who lost their jobs at the closed-down factories. The loose wires are therefore a painful reminder of the costs involved to those who really pushed for the outsourcing, thus motivating denials. These kind of denials are blameworthy, despite the psychological backstory. Saying that, there are also circumstances where organisational denials are clearly defensible, in the sense that objective outsiders would be likely to agree with the need for such measures. For example, national security concerns could provide such acceptable reasons for an organisational denial.

Coming to the level of individual members of an organisation for a moment, we might also not want to know. Maybe Bianca had an inclination as an experienced supply chain manager that something could be wrong, but did not seek to find answers. We might even actively avoid finding out about the consequences of our actions and choices. While I was working for an NGO that campaigned on ethical issues in global supply chains, an acquaintance once told me to never tell her anything bad about the multinational corporation she was working for, as she wanted to continue working for them. Although it was meant partly as a joke, this kind of attitude is typical of wilful ignorance, where there “is a self-interested reason for evading moral knowledge that might require one to rethink one’s way of life” (Isaacs 2011, p. 162). Corporations with market shares to protect can actively play into this, and thus we get meat packages with pictures of happy farm animals grazing on a green pasture, instead of pictures about the often bleak conditions under which animals are kept before being slaughtered.

The last category of ignorance arising from the suppression of knowledge are taboos. While taboos can be actively cultivated, leading to similar responsibility issues than those to do with secrecy and denials, I will suggest that taboos within collectives could also be created unintendedly through the pressure to converge. Szanto (2017) describes how reciprocal irrational influences can reinforce themselves both top-down and bottom-up in small groups and organisational and corporate contexts alike. The group members feel the bond between them, and strive for unanimity and in-group cohesion. These reciprocal expectations can sometimes override rational assessment of the best course for action, leading to what Janis (1982) has termed “groupthink”, where rationality and moral judgement deteriorates through in-group pressures.⁹ We could imagine something like this taking place among the

⁹ I thank Mikko Salmela for bringing the term to my attention. I find that groupthink could also be a contributing factor in an organisation resorting to denials.

top managers at Exciting Electronics, especially if there has been no previous product recalls and there was emphasis on not making mistakes. Therefore the option of a large-scale product call had become a kind of a taboo at Exciting Electronics, not seriously even entertained by the top managers.

Unintendedly created (i.e. non-cultivated) taboos mitigate blame to some degree at least, and perhaps in some cases they could even act as excusing condition. This does not, however, block forward-looking responsibility, as I will argue next.

4. Organisational design and the non-deliberate creation of ignorance

As I have been arguing, knowledge and ignorance are not evenly distributed within organizations. It is economical for an organization to consist of groups of experts that can work together when needed, as this allows for a wide range of skills and expertise to be employed in the organization. The process of specialization and coordination allows knowable recognised unknowns to be confined to or sustained within parts of the organization. When required, this ignorance can be overcome by bringing the different organizational actors together (Roberts 2013, p. 222). Ignorance, then, can be both necessary and even laudable, like with respecting privacy. However, organisational ignorance has problematic features also, as ignorance undermines the voluntariness and autonomy required for moral responsibility. In this section, I will argue that fragmentation of information in bureaucracies can lead to deliberate or non-deliberate creation of ignorance.

When we work together in collective settings, division of epistemic labour is not only very common, but also unavoidable. It allows organisations to absorb and process much greater amounts of facts than individual agents ever could. It also facilitates use of expert knowledge in overlapping areas and importantly the creation of new knowledge, ideas, inventions, and so on. We depend on others in our epistemic community for what Sandy Goldberg (2011, 121-122) calls *coverage* when we count on others to make relevant discoveries and reliably disseminate information of the same within the community.¹⁰ As with secrecy, I argue that the division of epistemic labour and fragmentation of information in organisations is acceptable as long as the reasons given for it would be salient for objective outsiders.

While it is natural that all knowledge is not shared, in some cases it can be hard to draw a clear line on what information should be available to whom. In bureaucratic organisation there is a requirement for some ignorance for the focus to be on roles, rather than personal characteristics. Roles are narrowly defined for a more or less precise purpose that serves the organisation's goals, therefore ignorance of

¹⁰ Thank you to an anonymous reviewer for bringing this term to my attention.

irrelevant personal characteristics of the people you deal with help things to run smoothly. An effective balance between informal relations and procedures, established in the course of frequent face-to-face contact, and the ignorance that is required for orderly procedures is thus required in any bureaucracy (Moore and Tumin 1949, p. 792-793). Here the ignorance is about facts that do not matter, so it is irrelevant and morally neutral.

However, organisations sometimes deprive individuals of their capacity to make good moral judgements by fragmenting available information. Recall the spy ring: maybe Shea would have refused to be part of the mission in the first place had she known details of the target and the method of the assignment. It could well be that the goal of the assignment goes against her values, and she was kept in the dark about the true nature of the mission on purpose to ensure her cooperation. To give a more humdrum example that is closer to real-life concerns, bureaucracy breaks work and knowledge into pieces, and bureaucratic compartmentalisation and the secrecy that often comes with it prevents information passing on from one department to another. This fragmentation of consciousness provides rationales for not knowing about problems, and for not trying to find out. Rational bureaucracy can, in this sense, stimulate irrationality (Jackall 1988, p. 194). Bureaucratisation is therefore never a purely technical matter, just a system of organisation, but a power system with privileges and domination. Max Weber already was worried about the implications of bureaucratisation for individuals' freedom and control, although he was supportive of bureaucracies as rational and efficient ways of humans to organise themselves.

Unlike Weber, Hannah Arendt (1970, pp. 38-39) was very critical of bureaucracies and described them as "rule by Nobody". Bureaucracies can compartmentalise work to such a degree that individual human action is reduced to mere behaviour. If division of labour goes too far, people no longer know what their role is in the larger organisation, what their work is linked to, what the results are. Responsibility is impossible to locate anymore and becomes so diffused that the people working in the bureaucracy can come to view their actions to be outside the normal human realm where they would be responsible for what they do. Expanding on Arendt's thoughts, Larry May (1996, pp. 71-76) similarly argues that organisational socialisation in bureaucracies can make people see themselves as the anonymous cogs of a machine, who do not have the need to develop a sense of responsibility in relation to what they do. Bureaucratic anonymity grows from the usual lack of face-to-face confrontation and not being directly linked to the consequences of one's actions. Some bureaucracies also socialise their members to feel that decisions should be made by the "experts" only, those members more experienced and knowledgeable. May (1996, p. 70) writes that "bureaucratic institutions socialize people to see

themselves not as actors but as those acted upon. The ensuing feelings of powerlessness can give rise to the acceptance of, and even participation in, harms these people [-] would never have found acceptable outside of the bureaucratic institution.”

In addition to fragmentation of information, organisational frameworks also affect the way we think. Our minds both organise and censor our experiences through conceptual schemes. Werhane (1999, pp. 85-95) describes how all of our activities are framed by mental models – our perspectives on things – and embedded in conceptual schemes. Our mental models are influenced by socialisation, culture, education, our upbringing, art, media, the place we work in. Our interests, desires, biases, intentions, and points of view operate as selective filters that restrict what we see in the world. Through the models, we make sense of our experiences, and interpret and clarify events to ourselves. This is often done retrospectively with events given a reframed focus and importance. We therefore do not observe the world objectively, but rather project our own perceptions on it and explain our experiences so that they fit our subjective point of view. We also tend to ignore data that does not fit our scheme. It is as if we are editing a movie and leave some of the scenes on the cutting room floor.

Thus corporate employees, for example, are trained to see things through the viewpoint of their employer, affecting the kinds of things they take into consideration when making decisions. The managers at Exciting Electronics could have been trained to think that they need to focus on introducing new product lines every quarter and keeping their prices as low as possible, for example. Maybe the ethos of Exciting Electronics gives low priority to proper quality control and safety measures in its supply chain, and places no emphasis whatsoever on the working conditions in its outsourced factories. If we choose any one perspective often, it gets reinforced in our minds. This is not to say that we have one-track minds, as most of us have several mental models to choose from so we can adapt to a given situation. Importantly, our perspectives can be altered if we choose to try to look at things from someone else’s perspective. In any case, mental models could lead to viewing certain information as unimportant and outside the focus of the organisation.

When it comes to mental models, especially those actively cultivated by the organisation, it is easy to see how we could argue that certain organisational practices are for example negligent and should be changed. An example could be a corporation assigning no importance in its internal practices on looking at the working conditions at its suppliers, e.g. by leaving such considerations off the check-lists it has created for its brand managers. The situation with regards to blame is less clear with the other mechanisms that produce ignorance in organisational settings. I have been arguing in this section that

fragmentation of information in bureaucracies can sometimes lead to non-deliberate creation of ignorance. When this has happened, and there is a harmful outcome, although not necessarily blameworthy for the outcome, the organisation should look at its design to try to make sure the same thing won't happen again. If they do not, I argue that they are blameworthy for being negligent.

5. Culpable ignorance

In the previous sections, I discussed how collectives can be culpable for the ignorance of their members. In this section, I turn to the possibility of culpable ignorance of individual members of the organisation (e.g. employees). In tracing cases of culpable ignorance (Smith 2016), an agent performs a morally inferior act from ignorance that can be traced back to an earlier act that created the conditions for ignorance. While this act of unwitting misconduct is excused by ignorance, the agent's earlier act is not, so the individual might be blameworthy. I will argue in this section that the earlier act that creates the conditions for ignorance can take place at the collective level, so blame should lie there also.

Assigning culpability to an agent for her ignorance is not yet to assign blameworthiness. Holly M. Smith (1983, p. 552) argues that all cases of culpable ignorance involve a sequence of acts: the initial act (the "benighting act") where the agent fails to improve their cognitive position, resulting in ignorance, and a subsequent act where the wrongful act is done due to this culpable ignorance. Smith further observes that frequently the benighting act takes the form of an omission, like failing to learn or find out something. The benighting act affects which subsequent acts are available to the agent, leading to the optimum act not being either epistemically or physically available to her.¹¹ According to Smith (2016), while the agent is not blameworthy for the act that was done in culpable ignorance, they are to blame for the earlier failure to obtain the information that would have led to her not being ignorant in the relevant manner. The agent has performed an act that is morally inferior to the counterfactual act she would have performed had she obtained all the necessary information. The ignorance is thus traceable to a past epistemic negligence.

To return to our example once again, it seems clear that Violet is not to blame in the tale of exploding toasters, as there is no way she could have known what was going to happen. She relied on what she heard from Exciting Electronics and in this way Violet was *epistemically dependent* on what Sharon told her about the products.¹² The case is not so clear with Sharon. Smith (1983, p. 544-547) presents three

¹¹ Alternative take on culpable ignorance is that culpability arises from holding of beliefs (Sher 2009).

¹² See Goldberg 2011 for a discussion on *epistemic dependence*.

types of cases where ignorance does not excuse, as the person should have realized what they were doing. *Deficient investigation* is the first type, either through failing to investigate properly, or failing to investigate at all. While this is not the case with Violet, Sharon could fall into this category if it were the case that she had failed to read an internal memo about possible problems with the new manufacturers' products. *Preventing subsequent discovery* presents the second case: a person has either failed to remove or introduced a condition which made it impossible for him to acquire true belief of x 's wrongness. If information about faulty products was available in print at Exciting Electronics but Sharon missed this because she has never learnt how to read properly (a life-long secret she has just about managed to hide from her employees), she would fall into this category. Finally, culpable ignorance could arise from *deficient inference*: had the agent made the inference warranted by his background beliefs, he would have correctly believed the act to be wrong. To use Sharon once again, had she remembered that her colleague Bob told her about problems in some new factory, she would have put two and two together when she heard from her manager Bianca that Exciting Electronics has changed its toaster manufacturer.

Although Smith's influential work on culpable ignorance looks at cases of individual responsibility, I see no reason why it cannot be framed in an organisational setting, allowing for the benighting act to be done by a different person than the morally inferior act that follows. Sharon did not disclose information to Violet about the faults, so Violet ordered faulty toasters to all shops in the retail chain. Had she known about the loose wires, she would have not completed the purchase. Violet's ignorance is not culpable here, though, while Sharon's might be. I gave examples earlier how Sharon's ignorance could be traced to either deficient investigation, deficient inference or preventing subsequent discovery. It could also be the case that Sharon's and Bianca's ignorance is not culpable either, but the result of suppression of information by Exciting Electronics, or a knowable recognised unknown for the organisation.

Compartmentalisation of information raises the further possibility that ignorance is produced systematically at Exciting Electronics, without it tracing back to the action or omission of any one person, or even a group of people. It could be several acts by several agents within the organisational setting that together produce the ignorance. Indeed, the deficient inference could take place at the collective level, with the organisation failing to make the inference warranted by their background knowledge and beliefs due to fragmentation of information, for example. In this way, the benighting act happens at the organisational level and any possible blame should be directed at the collective.

In an organisational setting the benighting act can either be done by an individual or take a collective form, in which case it is harder to point out exactly whose failure it was. Blameworthiness in a collective setting is a complex concept, so it will be harder to give simple or general answers.

Depending on how they are used and how justified their usage is in the first place, denials, secrecy, and taboos do not excuse everyone in an organisation, as they are usually instruments of power. They fall under suppression of knowledge at the top managerial level, or whatever level engages in the behaviour, while they can result in either excusing ignorance at the bottom level, or lead to conditions where it is all too easy to fall into the trap of making deficient inferences. If the parameters of a given role dictate acts that lead to culpable ignorance in others, the moral responsibility for that ignorance falls (either fully or at least in part) to the collective. The requirements of roles and an individual's leeway within them is a fascinating area for moral responsibility, but it falls outside the scope of this paper.

As I stated above, the benighting act can also take a collective form. I stipulated that the loose wires are so because of the too hectic manufacturing process. If competitiveness of the prices of goods takes precedent over all other considerations in their ethos, Exciting Electronics is taking a risk that falls within known risks. They might have been very lucky in the past and gotten away without adequate safety checks, but they could not justify this way of operating by appealing to ignorance: the knowledge would have been found out given the right resources and motivation. This would be an example of deficient investigation and the resulting ignorance is blameworthy.

Regardless, there is a potentially interesting consequence for work on individual culpable ignorance. I have argued that in an organisational context the benighting act, i.e. the earlier act that creates the conditions for ignorance, can take place at the collective level, so blame lies there also. If my argument works, then the quality of will of an agent is not necessary for blame, although it remains sufficient. The reason for this is that in an organisational setting there needs to be no bad will: organisational fragmentation of information alone can produce morally inferior acts. Bad will in this context should be understood as morally objectionable aversions or desires, or lack of proper moral concern. It often grounds the blameworthiness of agents in culpable ignorance literature. This common assumption is shared by Smith, who includes it in her account of moral blameworthiness.¹³ To give two other recent examples, Jan Willem Wieland (2017) suggests that individuals are blameworthy for their strategic ignorance depending on their moral concern, while Gunnar Björnsson (2017) conceptualises quality of

¹³ Smith 2016, p. 98: "I shall employ a 'quality of will' account according to which it is the quality of the agent's motivations in performing the blameworthy act that make her worthy of condemnation for performing it."

will of an agent through caring enough about how well things go and argues that ignorance fails to excuse when someone should have cared more.¹⁴

In an organisational setting there needs to be no bad will in order for there to be culpable ignorance. Violet and Sharon (and the workers in the factory where the toasters are put together) all lacked a morally objectionable configuration of aversions or desires. Sharon did not mislead Violet because she wanted to cause danger around the breakfast tables of Sell-A-Lot's toast-loving customers, or because she did not care enough about customer safety. It could very well be that she unwittingly lied about the quality and safety of the product because the relevant information was too fragmented at Exciting Electronics: division of labour between departments was too deep, lines of communication were unintentionally complicated, and so forth. Sharon relied on the epistemic coverage of her colleagues, but was let down in this regard. The act of telling the customers something that was not true was unwitting, as she would have acted otherwise had she been privy to the information about the faulty wires. However, from Violet's point of view Sharon's sale pitch about the high-quality of their toasters was a lie, taken as a statement representing the knowledge of Exciting Electronics.

Importantly, if the organisational design was unintentionally so that it caused fragmented information, there is no bad will at the level of the ethos of the organisation either. The benighting act takes a collective form, so the culpability is that of the collective also. We could of course come up with a version where the bad will of a manager in the factory that supplies Exciting Electronics with its machines, or a manager at Exciting Electronics or Sell-A-Lot, is the cause for the resulting organisational ignorance. This could be either due to personal failings on the part of the manager, or due to being too influenced by certain harmful organisational mental models. But this is beside the point: bad will is not necessary for the examples to get off the ground. Organisational fragmentation of information alone can produce morally inferior acts. Although there is no bad will, if the organisation does not take action to try to change its internal practices and communication flows after it has been made aware of the problems, then we can argue that it should care more. After all, it has come to know that the knowledge it has about the safety of its outsourced products is not robust enough to prevent such large-scale errors from happening. Still, some organisational fragmentation of information is always necessary and can lead to unintended consequences.

More generally, all collective action can result in outcomes that were completely unintended. Smith

¹⁴ The assumption is of course widely shared within moral responsibility literature, the quality of will account of blameworthiness is not exclusive to culpable ignorance literature.

argues that there is a degree to which we can be held accountable for the consequences of our actions, and this is linked to the outcomes falling within the predictable outcomes for that act. In other words, culpable ignorance arises only when the unwitting wrongful act falls within the known risks of the earlier act that infects the later act (Smith 1983, p. 551). This makes a separation between knowingly risking something and having no reason to believe that the benighting act would result in a wrongful act. Christopher Kutz's (2000) approach is somewhat different: while he acknowledges that our accountability for consequences that flow from our actions is in theory infinite and therefore needs to be "normatively delimited", he notes that some response is warranted even when the outcome is unforeseeable (pp. 142-143):

By taking responsibility for the consequences of our acts, we demonstrate to others a concern for their projects and interests, and thereby work to ensure their respect for our work. Within this delimited set of consequences, normative questions of individual response arise: whether to apologize, compensate, or repair.

Kutz illustrates this point with a wasp that enters the house when he lets the cat out. He did not intend to let the wasp in, but when it stings you badly, he should express sadness at your pain and offer you comfort. He is not at fault and resentment towards him is therefore not warranted. The response to the pain caused by the sting indicates the importance attached to your interests, and any claims for him to respond are rooted in the fact that his agency led to the suffering, however unintended it was. In other words, "accountability for unintended consequences manifest an acknowledgment of the fact that one's projects have interfered with another's interests" (p. 143).

This applies also to the unintended consequences of collective action, as we are complicit in the consequences of what we do together. If some harm is a direct consequence of what we do intentionally, even though the harm itself was not intended, we have a duty to acknowledge it in the appropriate way (i.e. apologise, compensate, etc.). If no apology or compensation was forthcoming from Sell-A-Lot to their customers, although the fault was with Exciting Electronics, the customers of the retail chain would be right to blame the company for not caring for their customers sufficiently.

But does it make sense to apply an account of reactive attitudes towards organisations in the first place? What, exactly, are we blaming when we are blaming a corporation, for example? Imagine an error done by a wholly automated organisation.¹⁵ Would liability work better here than assigning blame?

¹⁵ I thank Pekka Väyrynen for suggesting this option.

There are a number of possible routes to take to answer this question. One could point out that we do in fact blame collectives in our everyday lives and that it makes sense to blame the corporation because that is what we do. Reactive attitudes could be thought to be directed at the relevant members of the corporation, the ones who could make a difference, or towards the ethos. Or one could argue that reactive attitudes target the collective itself and that we should understand them in functional terms (like Hess and Björnsson 2017 do). What I want to argue here is that automation itself does not block out reactive attitudes.

Let us say that Exciting Electronics is ahead of its time, and is using an algorithm to make decisions about outsourcing its production. A complicated programme calculates the most cost-efficient supplier, with the best price-quality ratio, and sends out the necessary paperwork. The toasters start exploding. There was no option here for deficient inference in the way I described earlier with Sharon, Bob and Bianca: the algorithm alone knew about the switch in suppliers. Once it was informed about problems in the factory by the supplier (via an internet form), it calculated these to fall within acceptable risk parameters. Importantly, those parameters have been set by somebody. They are akin to engineering decisions that go into making automated vehicles. While the options might be set by engineers, the parameters of acceptable risks are decided by the corporation. Again, the ethos of Exciting Electronics could be argued to include negligence and to be blameworthy.

6. Concluding remarks

This paper has discussed how ignorance in an organisational setting is a complex phenomenon, and how it can be neutral, praiseworthy or blameworthy. Ignorance is a necessary ingredient to any organisation and it serves many purposes, from safeguarding trade secrets and our right to privacy, to ensuring the smooth running of daily operations. Ignorance is also a powerful tool for organisations to influence their members and stakeholders.

I have argued that the ethos of an organisation can be blamed for being morally lacking in some way, for example allowing negligence towards the safety of their customers or workers in their supply chain. While ignorance can be produced knowingly, it can also be an unintended side-effect of bureaucratization. If problems emerge, or a likelihood of bad outcomes is pointed out to them, the organisation has a forward-looking duty to try to fix its design. If the organisation fails to respond, they are blameworthy for failing to improve their design or organisational practices when it comes to

information flows.

One thing is certain: we should not ignore ignorance and its many facets when attempting to analyse organisations and other collective phenomena.

Acknowledgements

This paper has benefited from comments by the two anonymous reviewers and audience members at Philosophical Society of Finland Colloquium 2018. I would also like to thank the audiences at Collective Intentionality IX conference at Indiana University, Bloomington, and the subsequent invited talk at Marquette University, the 2014 Congress for Doctoral Students in Philosophy at University of Tampere, and the 2016 Stockholm-Helsinki graduate workshop in Practical Philosophy at Stockholm University for their very helpful comments on earlier versions of the argument. Last but not least, I would like to thank Arto Laitinen and Pekka Mäkelä for written comments on an early draft. The research was made possible by two personal grants from the Finnish Cultural Foundation, for which I am grateful.

Bibliography

- Arendt, Hannah (1970) *On Violence*. Harcourt Brace & Company, San Diego.
- Björnsson, Gunnar (2017) Explaining (Away) the Epistemic Condition on Moral Responsibility, in Robichaud and Wieland (eds): *Responsibility: The Epistemic Condition*. Oxford University Press, Oxford, pp. 146–162
- Björnsson, Gunnar and Kendy Hess (2017) Corporate Crocodile Tears?: On the Reactive Attitudes of Corporations. *Philosophy and Phenomenological Research* 94(2), pp. 273-298.
- Dufresne, Ronald L. and Evan H. Offstein (2008) On the Virtues of Secrecy in Organizations. *Journal of Management Inquiry* 17(2), pp. 102-106.
- Galison, Peter (2008) Removing Knowledge: The Logic of Modern Censorship, in Proctor and Schiebinger (eds) *Agnology: The Making and Unmaking of Ignorance*. Stanford University Press, Stanford, pp. 37-54.
- Goldberg, Sandy (2011) The Division of Epistemic Labor. *Episteme* 8(1), pp. 112-125.
- Hormio, Säde (2017) Can Corporations Have (Moral) Responsibility Regarding Climate Change Mitigation? *Ethics, Policy & Environment* 20(3), pp. 314–332.
- Isaacs, Tracy (2011) *Moral Responsibility in Collective Contexts*. Oxford University Press, New York.
- Jackall, Robert (1988) *Moral Mazes: The World of Corporate Managers*. Oxford University Press, Oxford.

- Janis, Irving (1982) *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*. Wadsworth, Boston.
- Kutz, Christopher (2000) *Complicity: Ethics and Law for a Collective Age*. Cambridge University Press, Cambridge.
- Laitinen, Arto (2014) 'Collective Intentionality and Recognition from Others', in Konzelmann Ziv and Schmid (eds) *Institutions, Emotions, and Group Agents: Contributions to Social Ontology*. Springer Science+Business Media Dordrecht, pp. 213-227.
- Lu, Lauren Xiaoyuan and Jayashankar M. Swaminathan (2015) Supply Chain Management. *International Encyclopedia of the Social & Behavioral Sciences (Second Edition)*, pp. 709–713, doi.org/10.1016/B978-0-08-097086-8.73032-7
- May, Larry (1996) *The Socially Responsive Self: Social Theory and Professional Ethics*. The University Chicago Press, Chicago.
- Moore, Wilbert E. and Melvin M. Tumin (1949) Some Social Functions of Ignorance. *American Sociological Review* 14(6), pp. 787-795.
- Peels, Rik (ed) (2016) *Perspectives on Ignorance from Moral and Social Philosophy*. Routledge, London.
- Proctor, Robert N. and Londa Schiebinger (eds) (2008) *Agnotology: The Making and Unmaking of Ignorance*. Stanford University Press, Stanford.
- Roberts, Joanne (2013) Organizational ignorance: Towards a managerial perspective on the unknown. *Management Learning* 44(3), pp. 215-236.
- Sher, George (2009) *Who Knew? Responsibility Without Awareness*. Oxford University Press, Oxford.
- Smith, Holly (1983) Culpable Ignorance. *The Philosophical Review* 92(4), pp. 543-571.
- Smith, Holly (2016) 'Tracing Cases of Culpable Ignorance', in Peels (ed.): *Perspectives on Ignorance from Moral and Social Philosophy*. Routledge, London, pp. 95-119.
- Szanto, Thomas (2017) Collaborative Irrationality, Akrasia, and Groupthink: Social Disruptions of Emotion Regulation. *Frontiers in Psychology* 7: 2002.
- Tuomela, Raimo (2007) *The Philosophy of Sociality: The Shared Point of View*. Oxford University Press, New York.
- Werhane, Patricia H. (1999) 'The Very Idea of a Conceptual Scheme', in Donaldson, Thomas; Werhane, Patricia H. & Cording, Margaret (eds) (2002): *Ethical Issues in Business: A Philosophical Approach (7th Edition)*. Pearson Education, New Jersey, pp. 83-97.
- Wieland, Jan Willem (2007) Responsibility for strategic ignorance. *Synthese* 194: pp. 4477-4497.